

Frederick J. Baker
1121 Via Del Rey
Goleta, California 93117

Chairman Ajit Pai
Commissioner Mignon Clyburn
Commissioner Michael O’Rielly
Federal Communications
Commission
445 12th Street, SW
Washington, DC 20554
RE: Docket No. 17-108, Restoring Internet Freedom

Dear Chairman and Commissioners:

Following, you will find my comments on the NPRM in question. In them, you will find that I both agree with and disagree with the NPRM’s assertions. My view is nuanced by my architectural view of the Internet and the services it delivers, and my understanding of the questions the Commission has wrestled with in the past 25 or more years.

/s/ Frederick J. Baker

1 [Introducing myself](#)

My name is Frederick Juergens Baker. I reside near Goleta, California, and am a US citizen. I speak as an individual.

I have been involved in the development and deployment of packet communications technologies since 1978, and specifically in IP communications since 1986. In 1989, I became involved with the [IETF](#), which is the premier standards development organization for that technology. In that context, I have written or edited sixty RFCs on a wide variety of topics. From 1996-2001 I served as the chair of the organization, and at this point serve as one of the co-chairs of its IPv6 Operations Working Group. I served in Focus Group 4 of NRIC V in 2001 representing Cisco Systems. In 2004, I participated in the FCC Technical Advisory Committee representing Cisco Systems. More recently, again at Cisco’s request, I have contributed to or edited several [BITAG](#) papers, some of which are referred to in the Open Internet Order. I am neither a lawyer nor an economist; however, I believe that I am qualified to comment on the technology in question.

In this document, I hope to respond to the statement “we seek comment”, which occurs repeatedly in the NPRM. I will organize my thoughts as they are organized in the NPRM, citing the paragraph to which I am responding.

2 Comments

2.1 Comment on the early paragraphs

A general comment that I would make on the recitation of history is that it is greatly abbreviated. As stated, there was a bipartisan consensus that the Internet should be open and free, and in 2010 the FCC decided it should not be. While those statements are true to a degree, to correctly understand the history there, I think the recitation needs to be enhanced to include the repeated discussions with law enforcement, which wanted some form of Lawfully Authorized Electronic Surveillance as described in Title II, and the issues of service delivery, service abuse, and anticompetitive behavior discussed under the rubric of Net Neutrality.

I do not personally agree with the FCC's decision to impose Title II on Internet Access, in large part because of the way it did it. The commission stated that a small set of provisions of Title II would be enforced, but most would not; some of their definitions and applications seem unduly strained. Selective enforcement has the unfortunate effect that the law can be effectively changed at any point without any of the visibility an NPRM such as this one has, by deciding that a previously unenforced provision needs to be enforced. If regulation is called for, I would rather see a regulation that was limited to what the commission intended to enforce, didn't have the other baggage, and was designed for the technology it was intended to regulate.

I did and do agree with the principles on which the Open Internet Order was based, as described in 2004-2010. If one wants a free and open Internet, a subscriber needs to be able to understand the service for which they have contracted, and know that communications they engage in and services they use are in fact the communications and services they intend. While a service provider certainly has the right to manage traffic in a way that facilitates quality service to all of its customers, it does not have the right to prevent lawful communications apart from a clear contractual provision. We call that the [End to End Principle](#), and by that mean that while the intervening network may determine for itself how best to serve its users, it should not do things that deny or subvert the intent of the user.

I would add one consideration to those principles, relating to privacy. We have a problem today with intelligence agencies and commercial entities eavesdropping on communications at will, based on fanciful interpretations of laws regarding lawful intercept, secret courts and court orders, routing of communications to come under obscure laws, and (in the case of commercial entities) the fact of handling the data. This has a demonstrable chilling effect on communication; Internet users develop what in the 1930's was called "the German look". I would argue that those who communicate should enjoy a reasonable expectation that their actual audience is a subset of their intended audience; this is part and parcel of anything resembling a right to privacy. When that reasonable expectation is violated, the event should be viewed as a violation, and governed by legal processes such as warrants, not an activity excused by the fact that government or commercial services can get away with it.

I also think that a change to the regulation that doesn't address the issues of Law Enforcement or anticompetitive behavior means that those issues will return and have to be dealt with again. I don't see the value of that.

One thing that bothers me in this NPRM is its title; that it is about "Restoring Internet Freedom". The issue is not, and has never been, about granting or denying freedom. It is about balancing freedoms. The user, who is purchasing a service but often has limited choices of whom to buy it from, should be able to make informed choices, and have his or her freedom guaranteed within the law. At the same time, the company providing that service needs to be free to innovate, and to prosecute its business in a cost-effective way. Any regulation that supports one without serving the other will eventually fail.

2.2 [Comments on paragraph 27](#)

This paragraph, which is the first to explicitly seek comment, would be greatly enhanced by a graphic showing the place of the broadband network's transmission protocol, the Internet Protocol, and application services. A service provider usually offers three classes of services: their favorite transmission technology (which might be cellular mobile, DOCSIS, DSL, GPON, WiFi, or one of several other things), end to end networking via IP, and various applications such as the web, mail, and so on. Their transmission technology basically conveys data from their users to their IP routers, and IP carries the information to systems somewhere else, whether in their network or another. The applications reside on servers, and are communication endpoints, not nodes in the path. In technical literature, this distinction is between the Data Plane, which contains transmission and IP, and the Control or Management Plane and applications, which together comprise everything else.

The NPRM reads, in part, "*We believe that Internet service providers offer the "capability for generating, acquiring, storing, transforming, processing, retrieving, utilizing, or making available information via telecommunications."*⁶⁸ *Whether posting on social media or drafting a blog, a broadband Internet user is able to generate and make available information online. Whether reading a newspaper's website or browsing the results from a search engine, a broadband Internet user is able to acquire and retrieve information online. Whether it's an address book or a grocery list, a broadband Internet user is able to store and utilize information online. Whether uploading filtered photographs or translating text into a foreign language, a broadband Internet user is able to transform and process information online. In short, broadband Internet access service appears to offer its users the "capability" to perform each and every one of the functions listed in the definition—and accordingly appears to be an information service by definition. We seek comment on this analysis. Can broadband Internet users indeed access these capabilities? Are there other capabilities that a broadband Internet user may receive with service? If broadband Internet access service does not afford one of the listed capabilities to users, what effect would that have on our statutory analysis? More fundamentally, we seek comment on how the Commission should assess whether a broadband provider is "offering" a capability. Should we assess this from the perspective of the user, from the provider, or through some other lens?*"

In the context, the paragraph doesn't describe the services of an ISP as much as it describes the services found on the Internet, some subset of which might also be offered by the ISP. Few service providers provide all of the examples given, but because they provide the data plane (transmission media and the IP layer), they provide access to the services that they provide and other services that are available from other providers. If the entirety of the Internet is kept in view, then yes, a user generally has access to those capabilities. If it is limited to the ISP's services, then no, a given ISP doesn't generally offer all of those services. I think it is fair to say that an ISP offers a service, such as the ability to exchange email, if it makes available to its subscribers a set of servers and names maintained by the ISP offering it. If it is web hosting, the ISP would need to provide a way that a user can upload a web page and have it available from an ISP-managed system, usually with a URL and domain name specified by the user. For email, it would entail the user being given an email address, and access to servers that can create or forward email *from* that address, and store and deliver messages that are *to* that address.

2.3 Comments on paragraph 28

Paragraph 28 wonders about the services that a consumer might use. While some services are certainly common, I don't think a list is helpful as it will be necessarily incomplete. In this context, the conceptual model in the previous section becomes very important. If the ISP offers an email service, a third-party email service is in competition with that service. The subscriber cannot communicate with either of them without the Data Plane, the transmission service and the IP-based Internet. However, either could be the email service the user uses. Personally, I use several different email services depending on context.

The ISP may have legitimate reasons to encourage its users to use its own service; for example, since a lot of web content is physically hosted in the United States, other parts of the world often find it in their interest to cache content, and to encourage their users to use the cache rather than consuming expensive bandwidth retrieving it from the source. However, this is exactly a case in which the intelligent user wants to know the provenance of the data; it may be stale (it may differ from what is currently available from the source), or may be corrupt, and it may have other problems. In addition, if the data is business-relevant (such as advertising), the source may have legitimate reasons to know that its data is not being cached or what has been delivered from the cache. So, it is inappropriate, in my view, to say that one model or another is absolutely right and another wrong. However, the user, as the party purchasing the service, has the absolute right, in my view, to know what service they are actually buying, and the advertiser may itself be a customer of a network.

The NPRM pointedly asks *"are we correct that offering Internet access is precisely what makes the service capable of "generating, acquiring, storing, transforming, processing, retrieving, utilizing, or making available information" to consumers?"* I would answer that sometimes the answer is "yes", and sometimes "no", and the question has to be interpreted in the context of a definition of the term "data". The transmission and IP

services, both of which are in the Data Plane, are absolutely required to deliver or to access a service, whether the subscriber is providing it, the ISP is, or a third party. If “data” is interpreted to mean the application payload of a transmission, such as a web page or portion thereof, then apart from the specific case of a transcoder, the data plane does not transform data, does not process it in any meaningful way, or “utilize” it, but the user has access to applications that do. If the “data” in view is information encoded in the IP header (including its IPv4 options or IPv6 extension headers) or a transport header such as TCP, then yes, specific configurations may process that data in specified ways, such as setting a flag to indicate that the packet experienced congestion en route, or changing a Differentiated Services Code Point to reflect the degree of conformance to a Service Level Agreement.

2.4 Comments on paragraph 29

The paragraph opens *“In contrast, Internet service providers do not appear to offer “telecommunications,” i.e., “the transmission, between or among points specified by the user, of information of the user’s choosing, without change in the form or content of the information as sent and received,” to their users. For one, broadband Internet users do not typically specify the “points” between and among which information is sent online. Instead, routing decisions are based on the architecture of the network, not on consumers’ instructions, and consumers are often unaware of where online content is stored.”*

I am reminded of the telephone in my home, growing up in the 1950’s. The black box had a dial on it, and a handset that I could hold to speak and to hear. If I dialed a number and the call was connected, I would find myself talking with someone via the instrument. I had no idea where they were, and I could not specify the routing of the call. If I called a number in use by a given store, for example, for all I knew the telephone I was calling could be temporarily located somewhere else or be answered by an employee of another store. The call might be routed directly from telephone to telephone, or circumnavigate the globe several times. I didn’t know, and if I had known, it was not under my control.

The logic of this opening sentence would suggest that the 1950’s telephone system that telephone was a part of was not a “telecommunication” system, because I could not specify the routing of the call. I think the people that ran it would have found that conclusion confusing.

On the other hand, Internet communications invariably requires the originator to specify the destination address (which is usually derived from a name such as a DNS name or a URL) the packet is intended to be delivered to. Additionally, both IPv4 and IPv6 include the optional capability to specify places along the way that the packet should visit; these are called “Source Route” options in IPv4, and the Type 0 Routing Header in IPv6, and in IPv6 there are other options that can affect routing. While these capabilities are defined in the specification, they are usually disabled to prevent certain kinds of attacks. The destination address may specify a specific system; in an anycast environment, it may also specify a service being offered by a set of computers that use

the same address. But if the user specifies that the intended destination of a packet is a given address, there is no doubt that it should not go somewhere else.

So, I would consider the statement to be only partially true. It is correct that Internet routing is determined independently from the user's direction; he doesn't say which sequence of data plane routers to use to get somewhere. It is also correct that Internet applications usually interpret a URL or DNS name to determine an IP address, which can be a complex process. However, the originator (or his computer) absolutely tells the system what the intended destination is, and routing seeks to implement that intent.

2.5 Comments on paragraph 30

What constitutes a "change in the form" of information?

I think the way the question is asked is important here, and if posed the question face to face, would find myself probing the interrogator.

Generally speaking, if a packet is sent using IPv4, it will follow an IPv4 route (which may be implemented over MPLS or other technologies, but that is a side point); if it is originated as IPv6, it will arrive as IPv6. That can change; in some networks, IPv4 packets are translated to IPv6 and delivered to the intended user, and in others, IPv6 packets are translated to IPv4 for delivery. I mentioned voice and video transcoding, which change the form of the data in the sense that it arrives encoded in a different way than it was sent. However, in the general case, the encoding (form) of the information does not materially change in flight. The packet a party receives is the one that it was sent, or would have been sent in the cases wherein translation or transcoding were required.

The paragraph asks, however, about the statistics of the data. In a DDOS attack, there might be a vast quantity of irrelevant traffic with a small quantity of useful communications hidden inside it. It is true that various technologies are used to separate between the two and focus on delivering the useful traffic. I would not describe that as a "change in form"; a change in form, to me, envisages changing the protocol, changing data encoding, or in some other way materially changing the useful data. This is about separating wheat from chaff, which to me is a very different question. The question of "relevance" brings up some concept of policy – a bank might consider accesses to critical financial systems to be among the "irrelevant" attack traffic if the attempt at access doesn't conform to its security requirements.

2.6 Comments on paragraph 33

I have chosen to not comment on paragraphs 31 or 32, as they seem questions of law more than of technology. I will say, though, in those paragraphs and in this one, as I said in section 2.1, the Commission seemed determined to make the law work, and made some uncomfortable choices. I would prefer that congress had written, and the FCC had implemented, the appropriate regulation rather than force fitting Title II.

2.7 Comments on paragraph 34

I spoke earlier about balancing freedoms. No, the Internet is not improved by reducing its security or increasing its complexity. In that, I would agree with the thought behind the paragraph. However, the Net Neutrality discussion started with blatantly anti-competitive behavior, and simply removing the regulations doesn't answer that concern.

2.8 Comments on paragraph 36

In 2.2, I distinguished between the data and control planes of service, and pointed out that (whether or not they actually are) they could be supplied independently. The first Internet Access contract I ever saw discussed the bit rate of the service, as did the one I most recently saw. However, at some point if one has enough bandwidth to do what one intends, other factors may tip the consumer's analysis in one direction or another. If I am looking for a web hosting service and have a choice of ISPs at essentially the same speed but only one of which offers the service, I might well choose that one as a result. So yes, the data plane (typified by a bit rate) and the application services are separate, but might be bundled by an ISP for reasons related to its business, and even if not bundled might be treated as a bundle by a potential customer.

2.9 Comments on paragraph 37

The 1983 Internet did not have DNS as a service; DNS was invented in 1984 as part of a solution to the absolute unwieldiness of an Internet without it. However, it is indeed an adjunct service; the router in my home uses OpenDNS rather than the DNS service offered by my carrier, and in Turkey, people have used Google's DNS service as opposed to Turkish Telecom's to circumvent censoring for a time.

However, nothing works without DNS today, whether it is the ISP's DNS server or someone else's. In 2.4, I said that the application supplies the address of a destination when a message is being sent. That address is almost invariably derived from DNS. No DNS, no destination address, no transmission, and no communication. The Internet just stops.

Caching (by which I presume you mean "of web content") is a very different story. Almost every web browser caches recently retrieved content, and in remote regions more centralized caching remains in use. However, it is a convenience; if content is not cached, it is retrieved, and the retrieval requires both time and capacity.

2.10 Comments on paragraph 38 and 39

The commission doesn't ask for comment here, but let me say that I agree that the Internet is an information service. If an information service is defined by the applications one can access, the Internet's applications engage in the exchange and change of information. The place where the distinction between Title I and Title II might be relevant relates to the data plane, the transmission media and the IP layer, which in context might correspond to an ATM Virtual Channel or similar transmission services in traditional telecom networks.

The NPRM asks, rhetorically, what harm has come from the “light touch” issue. I think law enforcement would respond that it has limited their investigatory capabilities, because they were not legally empowered to implement LAES. Part of the initial Net Neutrality discussion had to do with anticompetitive behavior by networks; if they are not common carriers as described in Title II, they are empowered to deliver services only to their direct customers, and might therefore limit accessibility by the customers of connected networks. As I said in 2.1, if these issues are not resolved – and I don’t see them addressed in this NPRM – the commission must expect them to recur and require resolution.

2.11 Comments on paragraph 40

The NPRM seeks comment on the analysis of the distinction between Title I and Title II services, and the validity of that distinction. I think that the transmission/IP services in the data plane have to be dealt with separately from the application services. The applications are without doubt information services. It may be valid to describe the transmission and IP services under Title II, or an appropriately redesigned variant of it that discusses only that which it enables and enforces. If the entirety of the service, including the data plane and the applications it supports access to are simultaneously in view, it must be treated as an information service.

3 Summary